

# TSM Copytool for Lustre HSM

Thomas Stibor

[t.stibor@gsi.de](mailto:t.stibor@gsi.de)

High Performance Computing  
GSI Helmholtz Centre for Heavy Ion Research  
Darmstadt, Germany

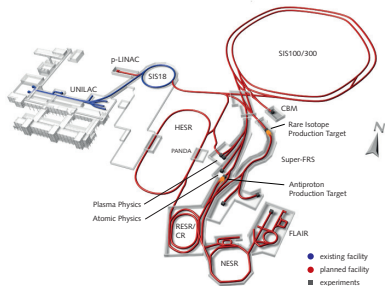
Tuesday 20<sup>th</sup> September, 2016

LAD 2016 Paris, France

# GSI/FAIR Overview

## FAIR: Facility for Antiproton and Ion Research

- Linear and Ring particle accelerators.
- Heavy Ion experiments.
- Medical irradiation facility for cancer therapy.



## Green IT Cube data center:

- Measures  $27 \times 30 \times 22$  meters, can hold 768 computer cabinets side by side on 6 floors.
- Highly energy-efficient, cooling with water.



# Lustre at GSI

## HPC at GSI/FAIR

- Green IT Cube data center
- Compute clusters
  - Prometheus (~9000 cores, QDR IB) [decommissioned]
  - Kronos (~10000 cores, FDR IB)
  - LCSC (~3200 cores + ~700 GPUs, FDR IB)
- Storage clusters
  - *Hera* (~7.3PB, Lustre 1.8.9 with Debian Squeeze 2.6 Kernel)
  - *Nyx* (~12PB, 70 OSSs, Lustre 2.5.3 on ZFS with Debian Wheezy 3.2 Kernel, Lustre 2.6.92 Clients on Debian Jessie 3.16 Kernel)
  - Currently in the process of moving data from *Hera*  $\Rightarrow$  *Nyx*, where Lustre client 1.8.9 mounts /hera (1.8.9 server) **and** /nyx (2.5.3 server).

GSI/FAIR is member of **Intel Parallel Computing Center** for developing a TSM Copytool for Lustre HSM.


# Overview of Lustre HSM and Copytools

Lustre with hierarchical storage management (HSM) feature is available since Lustre version 2.5 (partially landed in Lustre 2.4), since that manifold *copytools* are developed:

- Lustre Posix Copytool
- Lustre HPSS Copytool
- Lustre S3 Copytool
- Lustre Google Drive Copytool
- Lustre Droplet Copytool
- Lustre TSM Copytool (not yet released)


Lustre HSM  $\equiv$  seamlessly *archive*, *release* and *restore* data.

# Overview of HSM State Diagram




```
>dd if=/dev/zero of=zeros bs=1MiB count=32 conv=sync  
32+0 records in  
32+0 records out  
33554432 bytes (34 MB) copied, 0.401738 s, 83.5 MB/s
```

# Overview of HSM State Diagram



```
>lfs hsm_state ./zeros && ll -h zeros && du -h ./zeros
./zeros: (0x00000000)
-rw-r--r-- 1 root root 32M Sep  6 13:55 zeros
32M ./zeros
```


# Overview of HSM State Diagram



```
>lfs hsm_state ./zeros && ll -h zeros && du -h ./zeros
./zeros: (0x00000009) exists archived, archive_id:1
-rw-r--r-- 1 root root 32M Sep  6 13:55 zeros
32M ./zeros
```



# Overview of HSM State Diagram




```
>lfs hsm_state ./zeros && ll -h zeros && du -h ./zeros
./zeros: (0x0000000d) released exists archived, archive_id:1
-rw-r--r-- 1 root root 32M Sep  6 13:55 zeros
512 ./zeros
```






# Overview of HSM State Diagram



```
>lfs hsm_state ./zeros && ll -h zeros && du -h ./zeros
./zeros: (0x00000009) exists archived, archive_id:1
-rw-r--r-- 1 root root 32M Sep  6 13:55 zeros
32M ./zeros
```




# Overview of HSM State Diagram



```
>echo 0 >> ./zero && lfs hsm_state && ll -h zeros && du -h ./zeros
./zeros: (0x0000000b) exists dirty archived, archive_id:1
-rw-r--r-- 1 root root 33M Sep 19 09:16 ./zeros
32M ./zeros
```


# Overview of HSM State Diagram



```
>lfs hsm_archive ./zero && lfs hsm_state && ll -h zeros && du -h ./zeros
./zeros: (0x00000009) exists archived, archive_id:1
-rw-r--r-- 1 root root 33M Sep 19 09:16 ./zeros
33M ./zeros
```

## TSM Overview

Tivoli Storage Manager<sup>1</sup> (TSM) is a client/server software from IBM employed in heterogeneous distributed environments to *backup* and *archive* data.



<sup>1</sup>Now renamed to IBM Spectrum Protect.

# Tape Library



**Tape Library** is a storage device consists of

- tape drives,
- tape cartridges,
- barcode reader,
- tape robot

GSI employs *two* IBM 3584-L23 Tape Libraries with an overall capacity of  $1.2PB + 8.8PB$ .

# Backup vs Archive

*Backup:* A copy of the data is stored in the event the original becomes lost or damaged. Typically an incremental (forever) backup strategy is performed.

*Archive:* Remove from an on-line system those data no longer in day to day use, and place them into a long term retrievable storage (such as tape drives).

Lustre HSM is for *archiving* data.

# Some TSM Features

**Compression:** Compress data stream seamless either on client or server side.

---


**Deduplication:** Eliminating duplicate copies of repeating data.

---

**Collocation:** Store and pack data of a client in few number of tapes as much as possible to reduce the number of media mounts and for minimizing tape drive movements.

---

**Storage hierarchies:** Automatically move data from faster devices to slower devices based on characteristics such as file size or storage capacity.



Meta data is stored in a DB2 database (part of TSM server).

## Example Configuring Storage Hierarchies:

Example for setting data migration from small fast disk storage to slow large tape storage:

```
define devc fastdisks_devc devt=file maxcap=16G dir=/dir/dev/fastdisks
define devc slowtapes_devc devt=file maxcap=1024G dir=/dir/dev/slowtapes
define devc superslowtapes_devc devt=file maxcap=1048576G dir=/dir/dev/
    superslowtapes

define stg fastdisks_pool fastdisks_devc desc='fast disk storage pool' maxsize
    =15G nextstgpool=slowtapes_pool highmig=85 lowmig=40
define stg slowtapes_pool slowtapes_devc desc='slow tape storage pool' maxsize
    =1020G collocate=yes nextstgpool=superslowtapes_pool highmig=90 lowmig=50
define stg superslowtapes_pool superslowtapes_devc desc='super slow tape storage
    pool' maxsize=nolimit collocate=yes
```

If **high** threshold is reached, then move data to next storage pool hierarchy until first pool reaches the **low** threshold.

By means of *storage hierarchies* we can realized **storage caching layers**.



# TSM Installation Guide in KVM

Complete installation guide for settings up a TSM server e.g. within virtual machine (KVM) is provided at <http://web-docs.gsi.de/~tstibor/tsm/>.

TSM Server Installation Guide	
Table of Contents	
• 1. CentOS 7.2 and TSM Server 7.1.3	
◦ 1.1. Stop and disable Firewall	
◦ 1.2. Verify that firewall is disabled	
◦ 1.3. Disable SELinux	
◦ 1.4. Install the following packages	
• 2. Install TSM Server	
◦ 2.1. Download the latest TSM Server	
◦ 2.2. Console Installation	
◦ 2.3. Package Selection	
• 3. Login as user tsm and perform following steps	
◦ 3.1. Create directory <code>tsm-server</code> and database instance for user tsm	
◦ 3.2. Create <code>damsserv.opt</code> with following entries	
◦ 3.3. Initialize DB2 environment by adding following lines into <code>\$HOME/.profile</code>	
◦ 3.4. Create and format database and log locations	
◦ 3.5. Update file <code>userprofile</code> in <code>\$HOME/sqlib/userprofile</code>	
◦ 3.6. Install TSM Client for Testing	
◦ 3.7. Create file <code>/home/tsm/tsm-server/tsmdbmgr.opt</code> with content	
◦ 3.8. Create file <code>/opt/tivoli/tsm/client/api/bin64/dam.sys</code> with content	
◦ 3.9. Start TSM Server	
◦ 3.10. Create admin account and grant access	
◦ 3.11. Login from any machine with admin account as follows	
◦ 3.12. Register license	
◦ 3.13. Set Servname and IP Address	
◦ 3.14. Create storage directory and TSM device class assigned to this directory	
◦ 3.15. Create TSM storage group and assign to device class <code>tsdir_devic</code>	
◦ 3.16. Create TSM policyset and management class	
◦ 3.17. Create TSM copygroup for backup and archive	
◦ 3.18. Assign and activate TSM management class and policyset	
◦ 3.19. Register node and bind to domain	
• 4. Perform an archiving task	

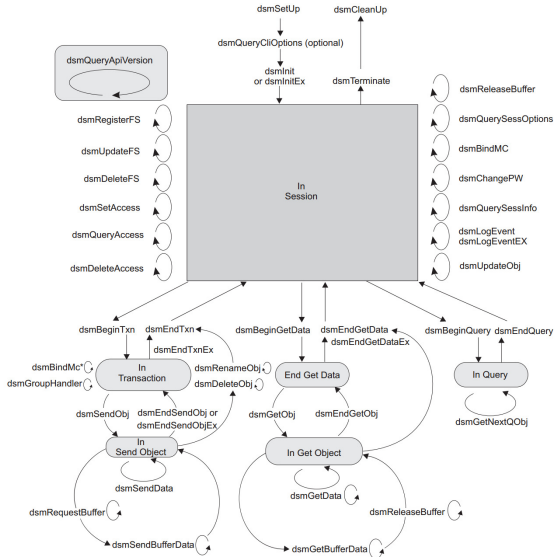
# TSM API Internals

Data on TSM side is stored in an object based format.

```
object # 1
fs: /, hl: /home/tstibor/dev/tsm/github-ltsm, ll: /README.md
object id (hi,lo) : (0,256004)
object info length : 32
object info size (hi,lo) : (0,12351)
object type : DSM_OBJ_FILE
object magic id : 71147
archive description : readme file description
owner :
insert date : 2016/5/23 16:13:54
expiration date : 2017/5/23 16:13:54
restore order (top,hi_hi,hi_lo,lo_hi,lo_lo): (722,0,16813,0,0)
estimated size (hi,lo) : (0,12351)

object # 2
fs: /, hl: /, ll: /tmp
object id (hi,lo) : (0,256003)
object info length : 32
object info size (hi,lo) : (0,86016)
object type : DSM_OBJ_DIRECTORY
object magic id : 71147
archive description : tmp directory description
owner :
insert date : 2016/5/23 16:13:37
expiration date : 2017/5/23 16:13:37
restore order (top,hi_hi,hi_lo,lo_hi,lo_lo): (722,0,16812,0,0)
estimated size (hi,lo) : (0,86016)
```

# TSM API Internals (cont.)



Taken from PDF document: Using the Application Programming Interface, Tivoli Storage Manager

# Initial TSMAPI (for interfacing Lustre HSM)

```
typedef struct {
    unsigned int magic;
    dsStruct64_t size;
    lustre_fid fid;
} obj_info_t;

typedef struct {
    char fpath[PATH_MAX + 1];
    char desc[DSM_MAX_DESCR_LENGTH + 1];
    obj_info_t obj_info;
    dsmObjName obj_name;
} archive_info_t;

dsInt16_t tsm_archive_file(const char *fs, const char *filename, const char *
    desc);
dsInt16_t tsm_archive_fid(const char *fs, const char *filename,
    const char *desc, const lustre_fid *fid);
dsInt16_t tsm_query_hl_ll(const char *fs, const char *hl, const char *ll, const
    char *desc, dsBool_t display);
dsInt16_t tsm_query_file(const char *fs, const char *filename, const char *desc,
    dsBool_t display);
dsInt16_t tsm_delete_file(const char *fs, const char *filename);
dsInt16_t tsm_delete_hl_ll(const char *fs, const char *hl, const char *ll);
dsInt16_t tsm_retrieve_file(const char *fs, const char *filename, const char *
    desc);
dsInt16_t tsm_retrieve_hl_ll(const char *fs, const char *hl, const char *ll,
    const char *desc);
```

# Console client `ltsmc` based on TSM API

Simple console client for testing and demonstrating

- `tsm_archive_file(const char *fs, const char *filename, const char *desc);`
- `tsm_archive_dir(const char *fs, const char *directory, const char *desc);`
- `tsm_query_hl_ll(const char *fs, const char *hl, const char *ll, const char *desc, dsBool_t display);`
- `tsm_query_file(const char *fs, const char *filename, const char *desc, dsBool_t display);`
- ...

```
> Syntax: bin/ltsmc
-a, --archive
-r, --retrieve
-q, --query
-d, --delete
-f, --fsname <STRING>
-h, --hl <STRING>
-l, --ll <STRING>
-c, --description <STRING>
-n, --node <STRING>
-u, --username <STRING>
-p, --password <STRING>
-s, --servername <STRING>
-v, --verbose (optional level <v,vv,vvv>)
```

# Live Demo

```
BI | tsm@tsm-centos-7-2 ~
option -l, --ll/*
[DEBAG] [src/tsnapi.c] (dsnInitEx:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnRegisterFS:handle:1:ANS0242M (RC2862) On dsnRegisterFS the filepath is already registered)
[DEBAG] [src/tsnapi.c] (dsnQuerySessOptions:handle:1:ANS03021 (RC0) Successfully done.)
[VERBOSE]
DSMT_DIR : /opt/tivoli/tsm/client/api/bin64
DSMT_CONFIG : /home/tstibor/dev/tsm/tsm/dsmopt/dsm.opt
serverName : LXDVB1-KVM-TSM-SERVER
compressMethod : 1
serverAddress : 192.168.254.101
nodeName : LXDVB1
compress : 0
compressalways : 1
passwordAccess : 0
[DEBAG] [src/tsnapi.c] (dsnQuerySessInfo:handle:1:ANS03021 (RC0) Successfully done.)
[VERBOSE]
Server's ver.rel.leve : 7.1.3.0
ArchiveRetentionProtection : No
[VERBOSE]
Max number of multiple objects per transaction: 4096
Max number of Bytes per transaction: 26214400
dsnSessInfo.tideline: /
dsnSessInfo.hideline: /
[VERBOSE]
API Library Version = 7.1.3.0
[VERBOSE]
tsm query archive with settings
fs: /
hl: /home/tstibor/dev/tsm/tsm/archives/letters
ll: /*
owner: *
descr: *
[DEBAG] [src/tsnapi.c] (dsnBeginQuery:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnGetNextObj:handle:1:ANS02581 (RC2200) On dsnGetNextObj or dsnGetData there is more available data)
[DEBAG] [src/tsnapi.c] (dsnGetNextObj:handle:1:ANS02581 (RC2200) On dsnGetNextObj or dsnGetData there is more available data)
[DEBAG] [src/tsnapi.c] (dsnGetNextObj:handle:1:ANS02581 (RC2200) On dsnGetNextObj or dsnGetData there is more available data)
[DEBAG] [src/tsnapi.c] (dsnGetNextObj:handle:1:ANS02721 (RC121) The operation is finished)
[DEBAG] [src/tsnapi.c] (dsnBeginTxn:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnDeleteObj:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnEndTxn:handle:1:ANS03021 (RC0) Successfully done.)
[VERBOSE]
Deleted obj fs: /
hl: /home/tstibor/dev/tsm/tsm/archives/letters
ll: /a.big.pdf
[DEBAG] [src/tsnapi.c] (dsnBeginTxn:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnDeleteObj:handle:1:ANS03021 (RC0) Successfully done.)
[DEBAG] [src/tsnapi.c] (dsnEndTxn:handle:1:ANS03021 (RC0) Successfully done.)
[VERBOSE]
Deleted obj fs: /
hl: /home/tstibor/dev/tsm/tsm/archives/letters
ll: /a.big.pdf
+ set xx
##### deleting h/ll done #####
Successfully archived and retrieved data verified with MD5SUM
tstibor@lxdvb1:~/dev/tsm/tsm$
```

```
BI | jnar/bn/ush
name:
bin/ltsmc --archive --fsname '/' -c 'Complete Historic Linux Kernel' --node lxdvb1 --password lxdvb1 /t
mp/archives/linux
### Querying Data
For querying all data stored in directory '/tmp/archives/linux'
bin/ltsmc --query --fsname '/' --node lxdvb1 --password lxdvb1 --hl '/tmp/archives/linux/*' --ll '/'
Querying the single file '/tmp/archives/linux/Makefile'
bin/ltsmc --query --fsname '/' --node lxdvb1 --password lxdvb1 --hl '/tmp/archives/linux' --ll '/Makefi
le'
### Retrieving Data
We first delete all data in '/tmp/archives/linux' and the restore the data to that directory:
rm -rf /tmp/archives/linux && bin/ltsmc --retrieve --fsname '/' --node lxdvb1 --password lxdvb1 --hl '/'
tmp/archives/linux/* --ll '/'
### Deleting Data
To delete data on the TSM perform the following command:
bin/ltsmc --delete --fsname '/' --node lxdvb1 --password lxdvb1 --hl '/tmp/archives/linux/*' --ll '/'
A subsequent "delete" or "query" command shows that there is no data left on the TSM side.
bin/ltsmc --delete --fsname '/' --node lxdvb1 --password lxdvb1 --hl '/tmp/archives/linux/*' --ll '/'
### TODOs
* Integrate LTSNAPI into the Lustré HSM framework.
E|
```

```
BI | tsm@tsm-centos-7-2 ~ tsm-server
ANR04061 Session 358 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49791)).
ANR04031 Session 358 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 359 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49792)).
ANR04031 Session 359 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 360 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49793)).
ANR04031 Session 360 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 361 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49794)).
ANR03401 FILE volume /home/tsm/tsm-storage/0000000C.BFS mounted.
ANR05111 Session 361 opened output volume /home/tsm/tsm-storage/0000000C.BFS.
ANR05141 Session 361 closed volume /home/tsm/tsm-storage/0000000C.BFS.
ANR04031 Session 361 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 362 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49795)).
ANR04031 Session 362 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 363 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49796)).
ANR03401 FILE volume /home/tsm/tsm-storage/0000000C.BFS mounted.
ANR05101 Session 363 opened input volume /home/tsm/tsm-storage/0000000C.BFS.
ANR05141 Session 363 closed volume /home/tsm/tsm-storage/0000000C.BFS.
ANR04031 Session 363 ended for node LXDVB1 (GNU/Linux).
ANR04061 Session 364 started for node LXDVB1 (GNU/Linux) (Tcp/Ip 192.168.254.254(49797)).
ANR04031 Session 364 ended for node LXDVB1 (GNU/Linux).
ANR3205W A full database backup will be started. The archive log space used is 85% and the archive
log space used threshold is 80%.
ANR2994E The automatic database backup was terminated.
ANR3205W A full database backup will be started. The archive log space used is 85% and the archive
log space used threshold is 80%.
ANR2994E The automatic database backup was terminated.
```

Demo available

at <http://web-docs.gsi.de/~tstibor/tsm/tsm-screen-cast-1.mp4>

## Summary & Outlook

- Currently in the process of hooking TSM API into `llapi_hsm*` functions for finalizing TSM copytool.
- TSM API and `ltsmc` is released <https://github.com/tstibor/ltsm>, if Lustre file system is decommissioned, then `ltsm` still can be used to restore data.

Thank you & questions